



US009277480B2

(12) **United States Patent Kohli**

(10) **Patent No.:** US 9,277,480 B2
(45) **Date of Patent:** Mar. 1, 2016

(54) **CLOUD CONTROLLER FOR NEXT GENERATION DATA NETWORK**

H04W 40/34 (2013.01); *H04W 56/00* (2013.01); *H04W 72/0453* (2013.01); *H04L 1/1867* (2013.01); *H04W 72/00* (2013.01)

(71) Applicant: **Facebook, Inc.**, Menlo Park, CA (US)

(58) **Field of Classification Search**
None
See application file for complete search history.

(72) Inventor: **Sanjai Kohli**, Manhattan Beach, CA (US)

(73) Assignee: **FACEBOOK, INC.**, Menlo Park, CA (US)

(56) **References Cited**

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 140 days.

U.S. PATENT DOCUMENTS

(21) Appl. No.: **14/214,956**

9,215,644	B2	12/2015	Kohli	
2005/0283643	A1 *	12/2005	Banerjee	H04L 41/0668 714/4.3
2006/0114818	A1 *	6/2006	Canali	H04L 45/00 370/216
2006/0146696	A1 *	7/2006	Li	H04L 12/5695 370/218
2008/0069034	A1 *	3/2008	Buddhikot	H04W 40/14 370/328
2008/0248802	A1	10/2008	Krishnamoorthy	
2011/0038253	A1 *	2/2011	Yabusaki	H04L 45/02 370/217
2011/0045767	A1	2/2011	Rofougaran et al.	
2011/0159821	A1	6/2011	Park	

(22) Filed: **Mar. 16, 2014**

(65) **Prior Publication Data**

US 2014/0286251 A1 Sep. 25, 2014

Related U.S. Application Data

(60) Provisional application No. 61/789,188, filed on May 15, 2013.

(Continued)

OTHER PUBLICATIONS

U.S. Appl. No. 14/214,957 by Kohli, S., filed Mar. 16, 2014.

(Continued)

(51) **Int. Cl.**

<i>H04W 40/14</i>	(2009.01)
<i>H04W 40/02</i>	(2009.01)
<i>H04W 40/12</i>	(2009.01)
<i>H04W 24/04</i>	(2009.01)
<i>H04B 15/00</i>	(2006.01)
<i>H04L 12/707</i>	(2013.01)
<i>H04W 40/34</i>	(2009.01)
<i>H04W 72/04</i>	(2009.01)
<i>H04W 28/18</i>	(2009.01)
<i>H04W 56/00</i>	(2009.01)
<i>H04W 72/00</i>	(2009.01)
<i>H04L 1/18</i>	(2006.01)

Primary Examiner — Benjamin H Elliott, IV
(74) *Attorney, Agent, or Firm* — Perkins Coie LLP

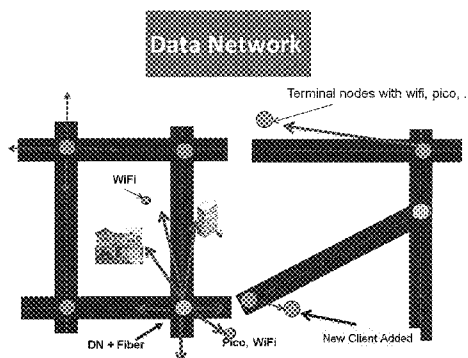
(52) **U.S. Cl.**

CPC *H04W 40/14* (2013.01); *H04B 15/00* (2013.01); *H04L 45/22* (2013.01); *H04W 24/04* (2013.01); *H04W 28/18* (2013.01); *H04W 40/02* (2013.01); *H04W 40/12* (2013.01);

(57) **ABSTRACT**

A Next Generation Data Network is described. It leverages the “cloud” for data management, low frequency data computation and analytics. The wireless network is a single frequency network that permits limited non-line of sight operation. The wireless network using packet switched beams, the beams are formed and switched electronically. It utilizes advanced signal processing to compensate for low transmit signal power and multipath reflections that can be frequency or flat fades.

40 Claims, 19 Drawing Sheets



Network capacity increases 2X for each fiber point of presence visioning

(56)

References Cited

OTHER PUBLICATIONS

U.S. PATENT DOCUMENTS

2013/0128726	A1	5/2013	Hellhake et al.	
2013/0201857	A1	8/2013	Bhargava et al.	
2013/0301619	A1	11/2013	Singh et al.	
2014/0074392	A1*	3/2014	Holm	G01C 21/3415 701/400
2014/0286156	A1	9/2014	Kohli et al.	
2014/0286251	A1*	9/2014	Kohli	H04B 15/00 370/329

Notice of Allowance mailed Sep. 29, 2015, for U.S. Appl. No. 14/214,957 of Kohli, S., filed Mar. 16, 2014.

U.S. Appl. No. 14/878,943 by Kohli, S., filed Oct. 8, 2015.

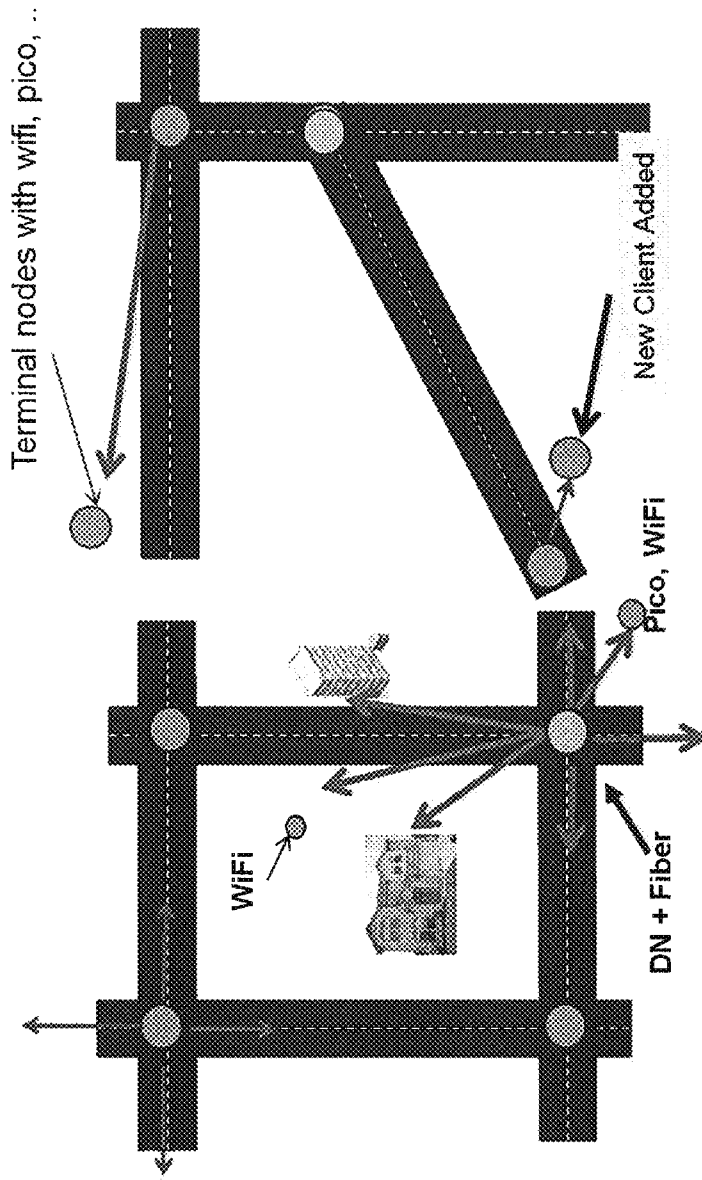
U.S. Appl. No. 14/878,960 by Kohli, S., filed Oct. 8, 2015.

U.S. Appl. No. 14/878,974 by Kohli, S., filed Oct. 8, 2015.

U.S. Appl. No. 14/878,982 by Kohli, S., filed Oct. 8, 2015.

* cited by examiner

Figure 1: Data Network



Network capacity increases 2X for each fiber point of presence visioning

Figure 2: Next Generation Data Network

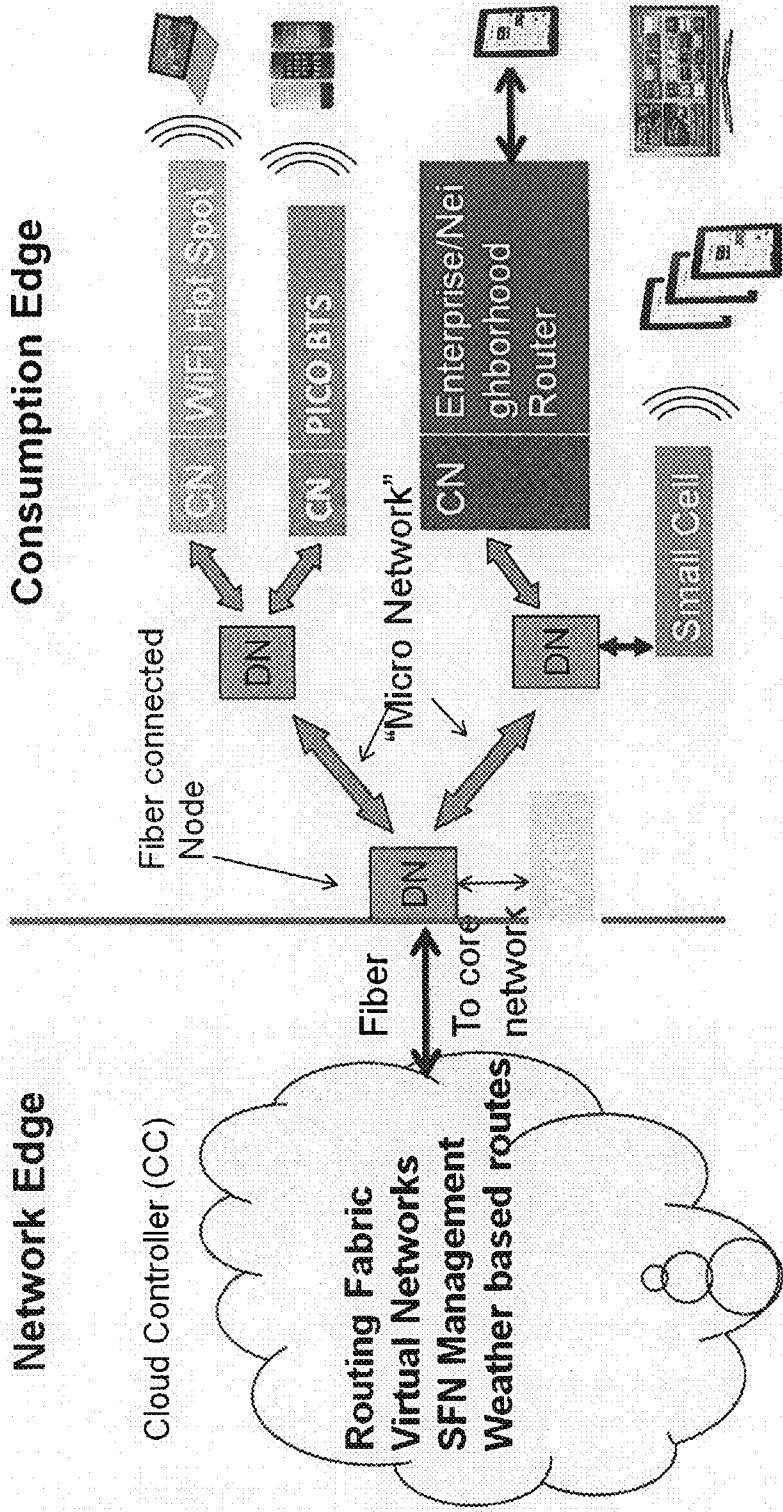
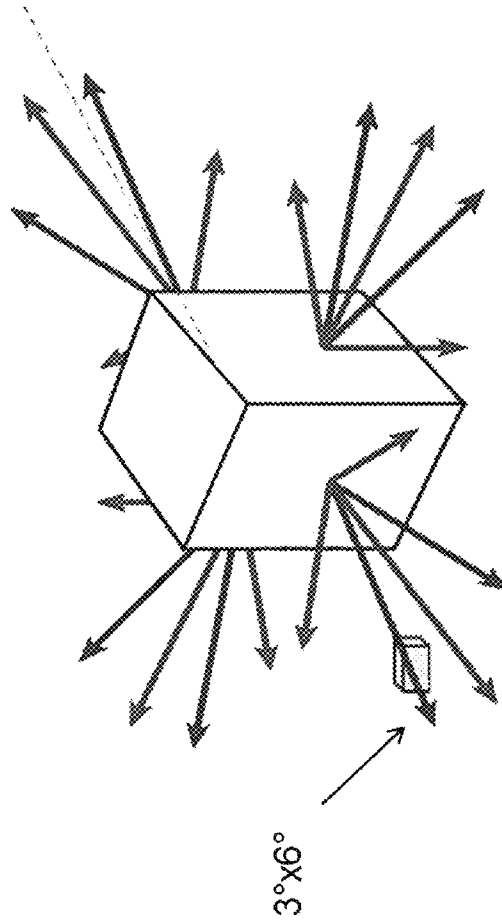


Figure 3: Distribution Node



360° coverage via
4 synchronous beams

Figure 4: SubNets & Fiber redundancy

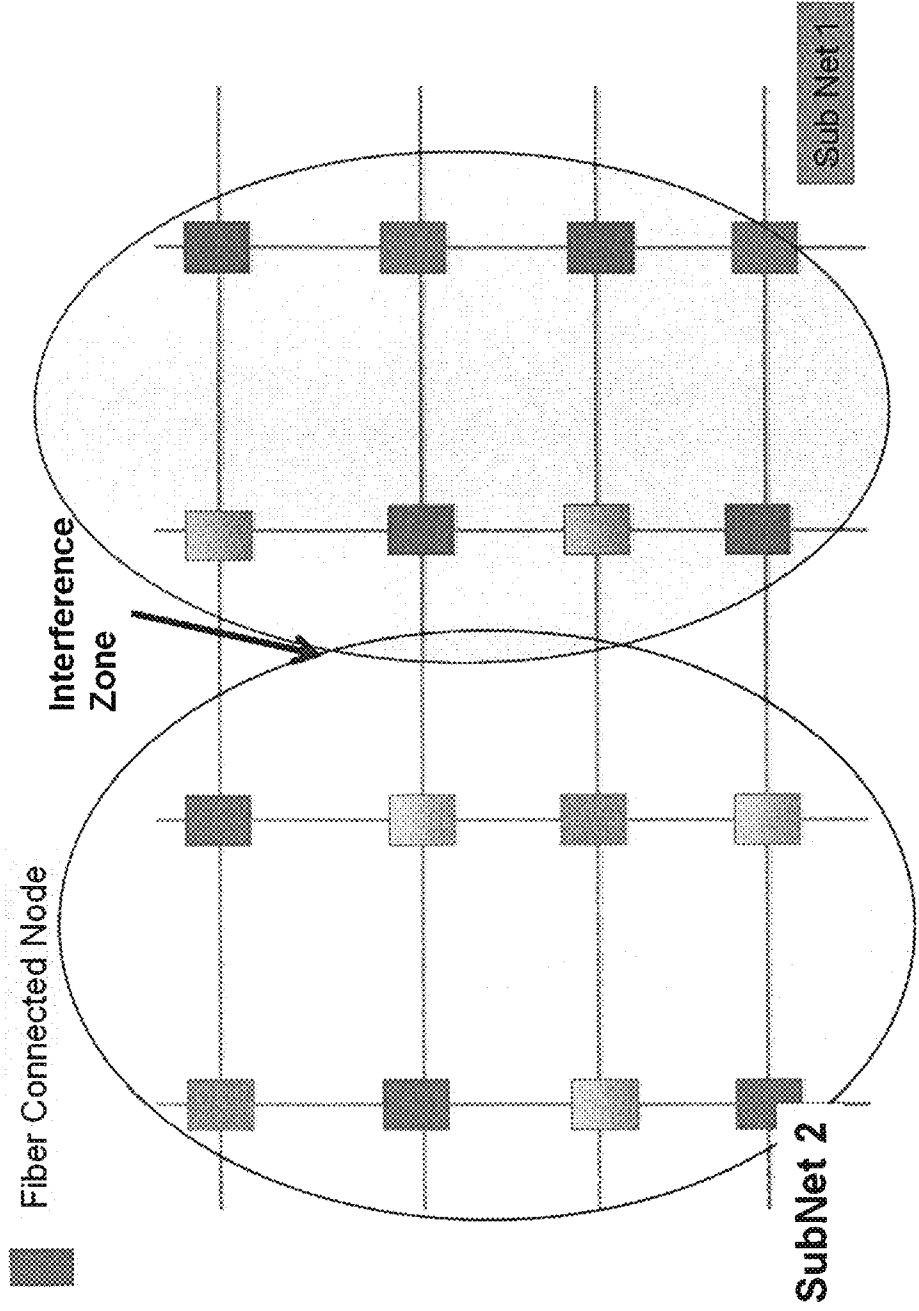
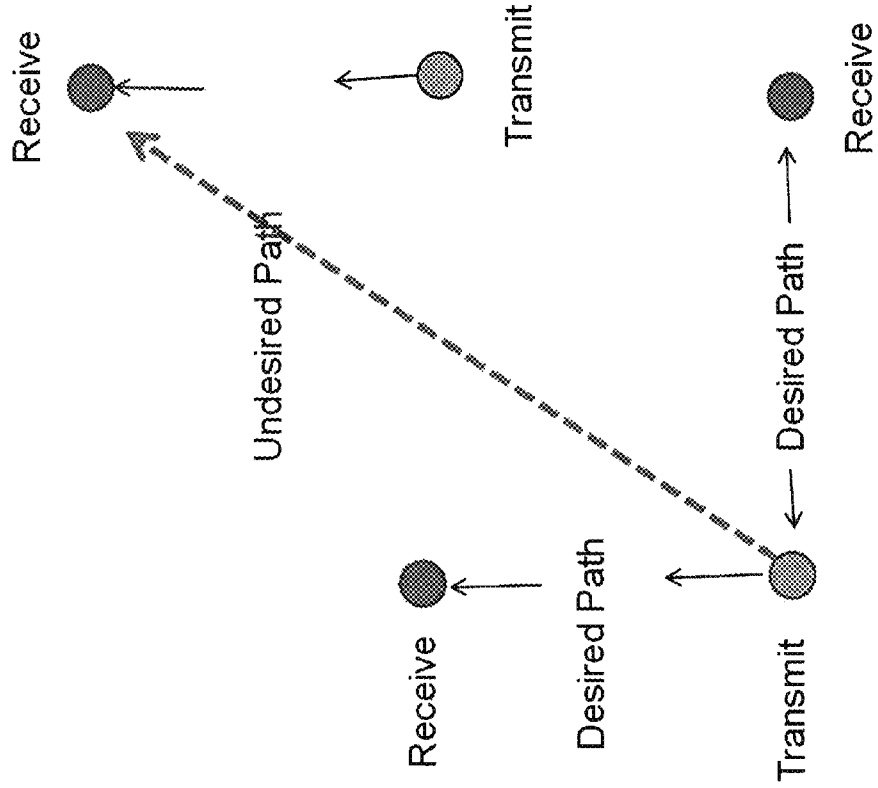
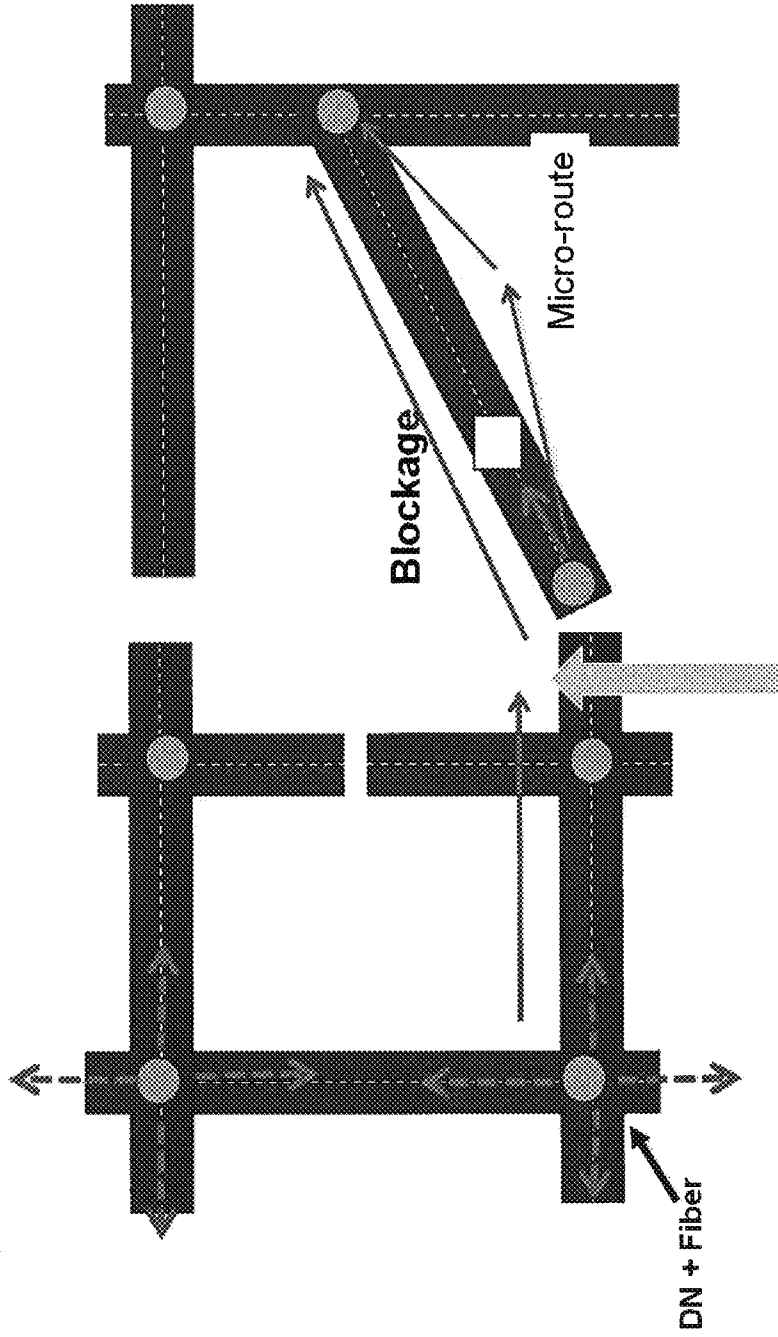


Figure 5: RF Maps Characterize All Undesired Paths in the Network

- ❑ Exception Pairs where Transmit Receive cannot be performed simultaneously
- ❑ Defined as a set of transmit receive pairs
- ❑ Beam parameters
- ❑ Associated Mac addresses for source and destination packets



**Figure 6:
Micro Route Around a Physical or Interference Blockage**



Original Cloud Computed Macro Routes

Figure 7:
Re-Route Around a Physical or Interference Blockage

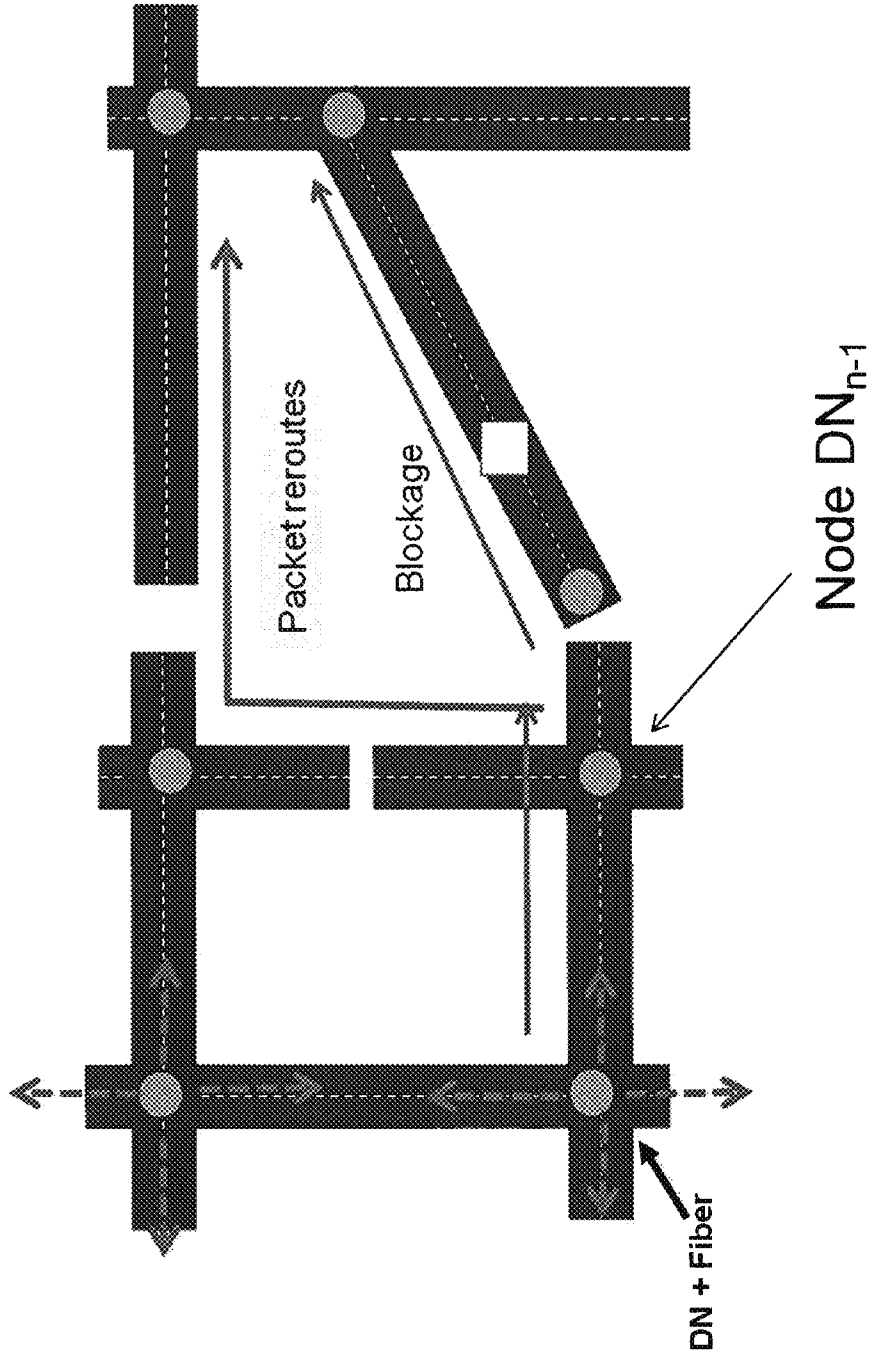
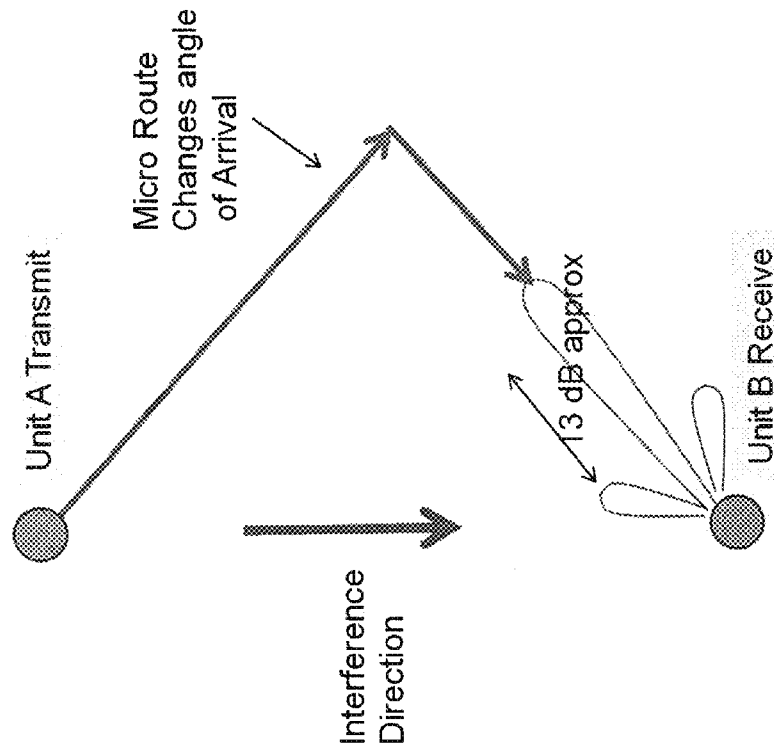


Figure 8: Interference Suppression using a Micro Route



Inovi proprietary & confidential

Figure 9: Spatial Interference Canceller

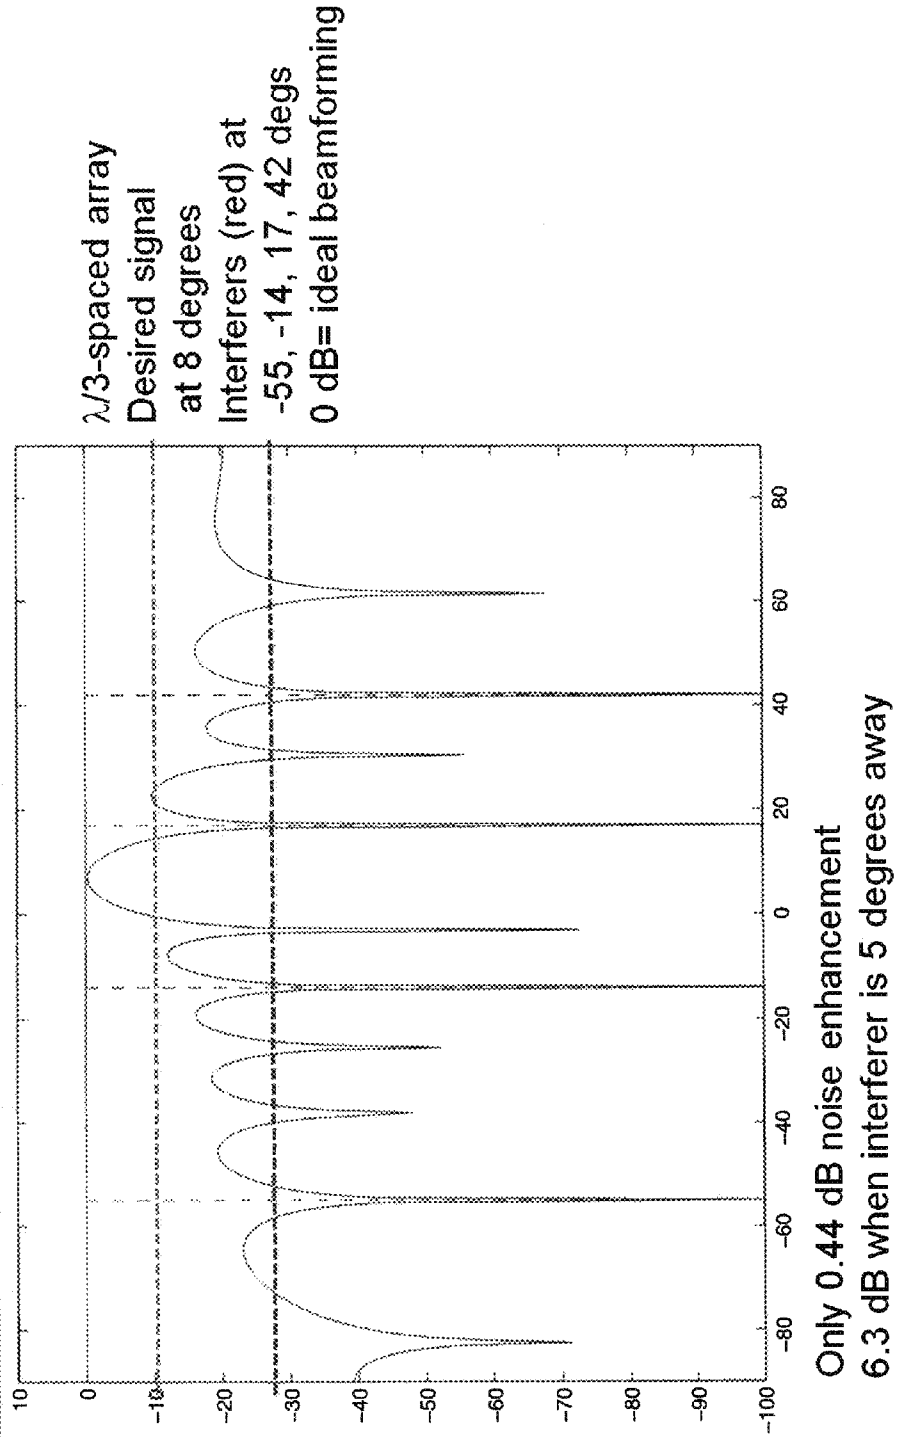
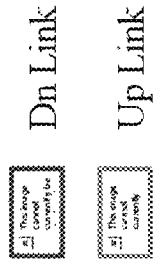
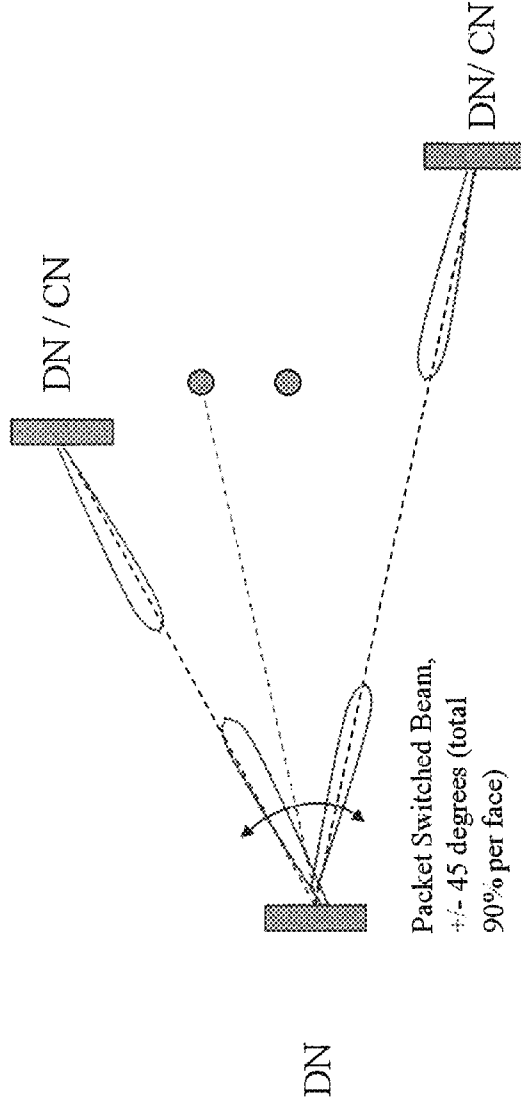


Figure 10: Packet Switched Beams For Coverage & Range



Adaptive TDMA-TDD structure



Each sector has a single Tx (or Rx) beam time multiplexed with Up/Dn Duplexing

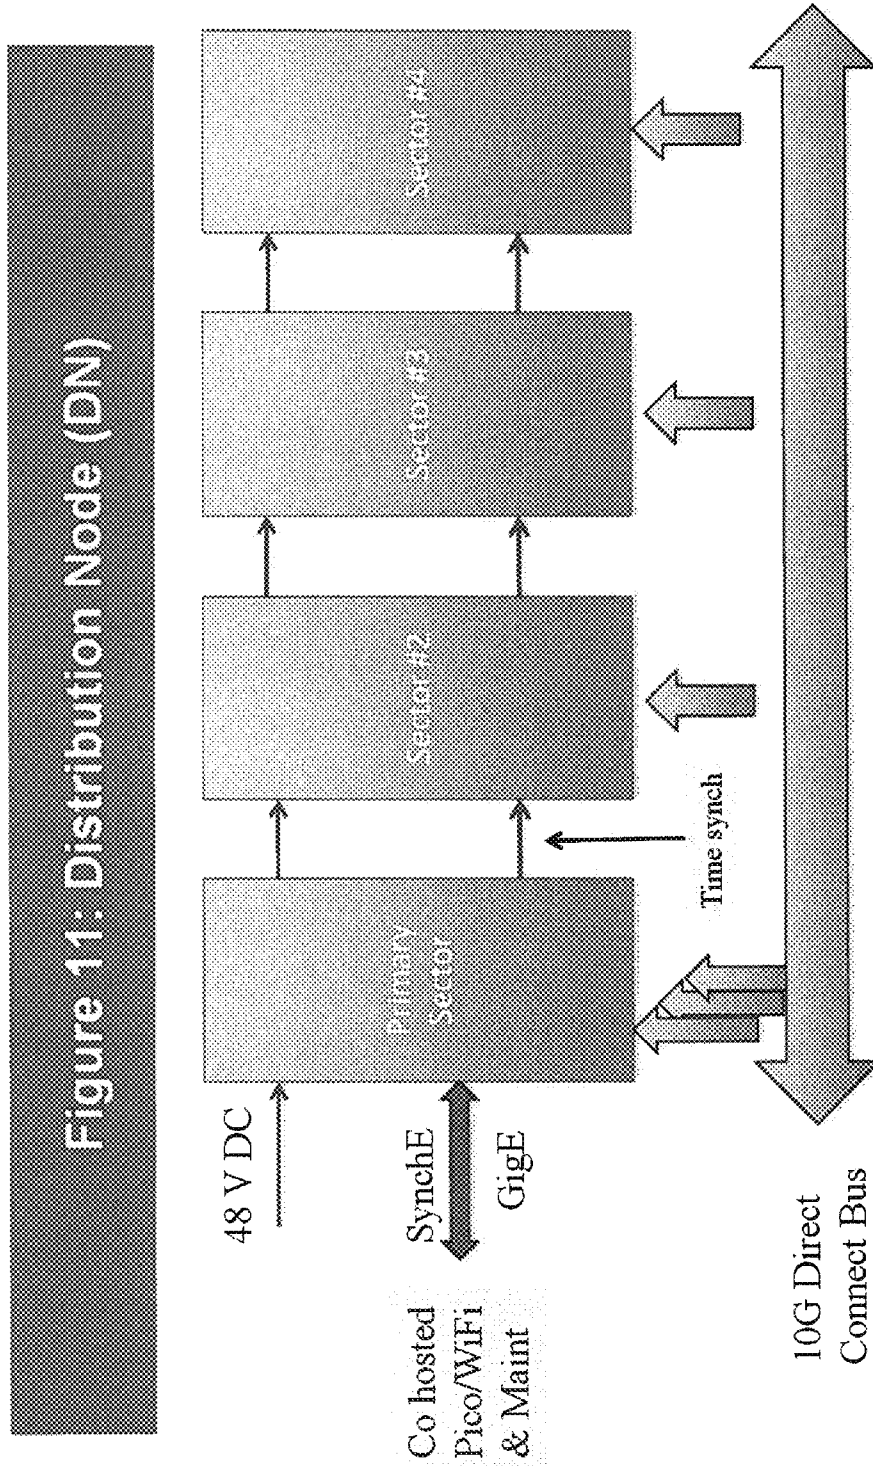


Figure 12: Using Side lobes for transmission

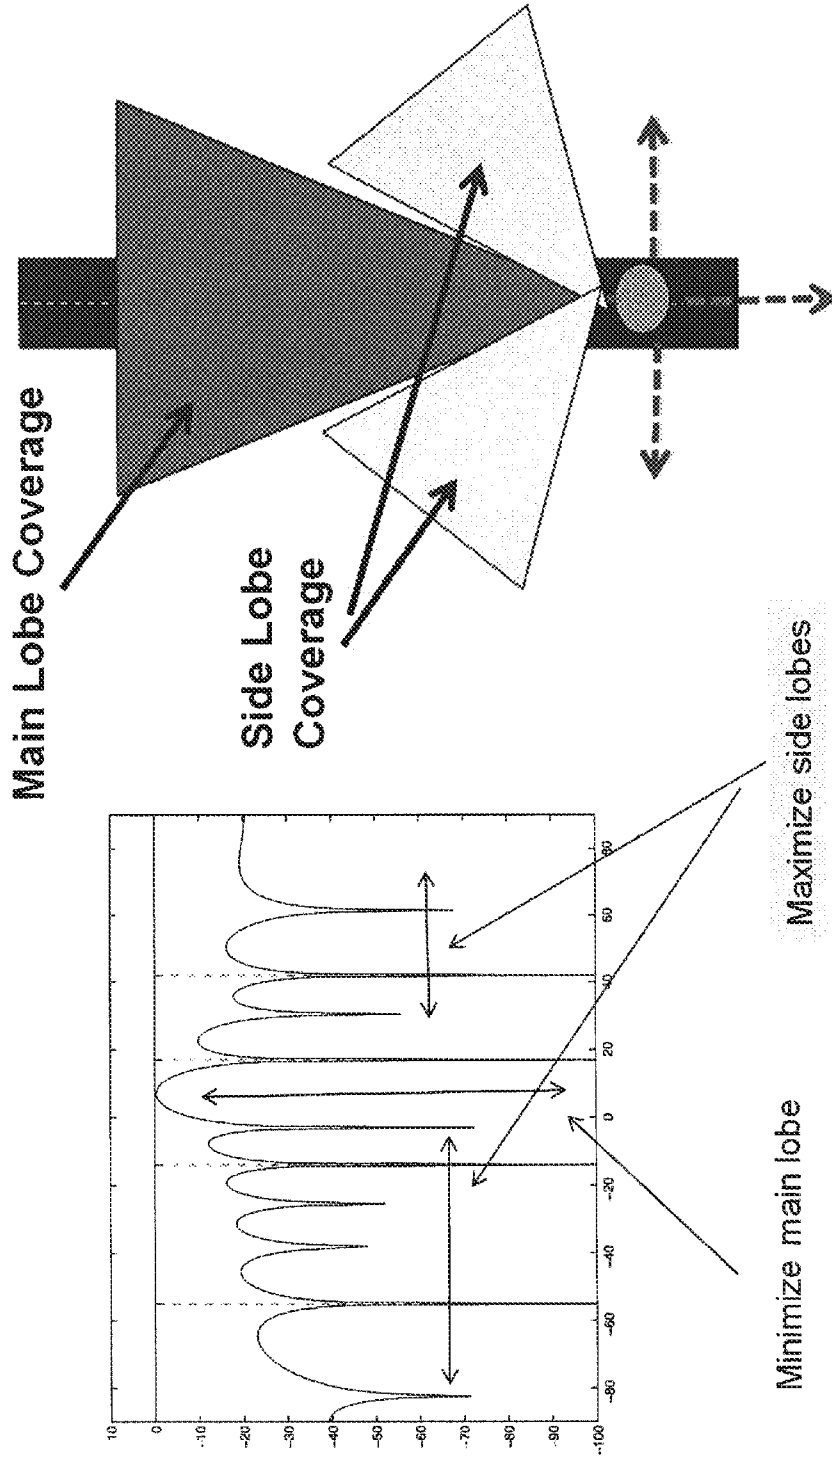
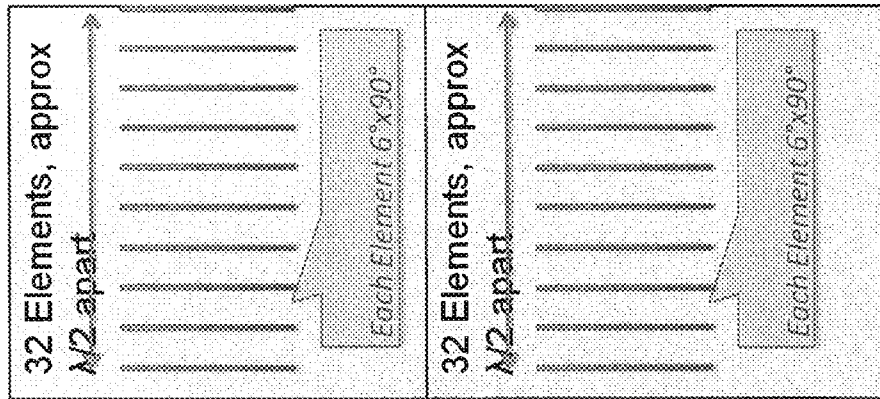
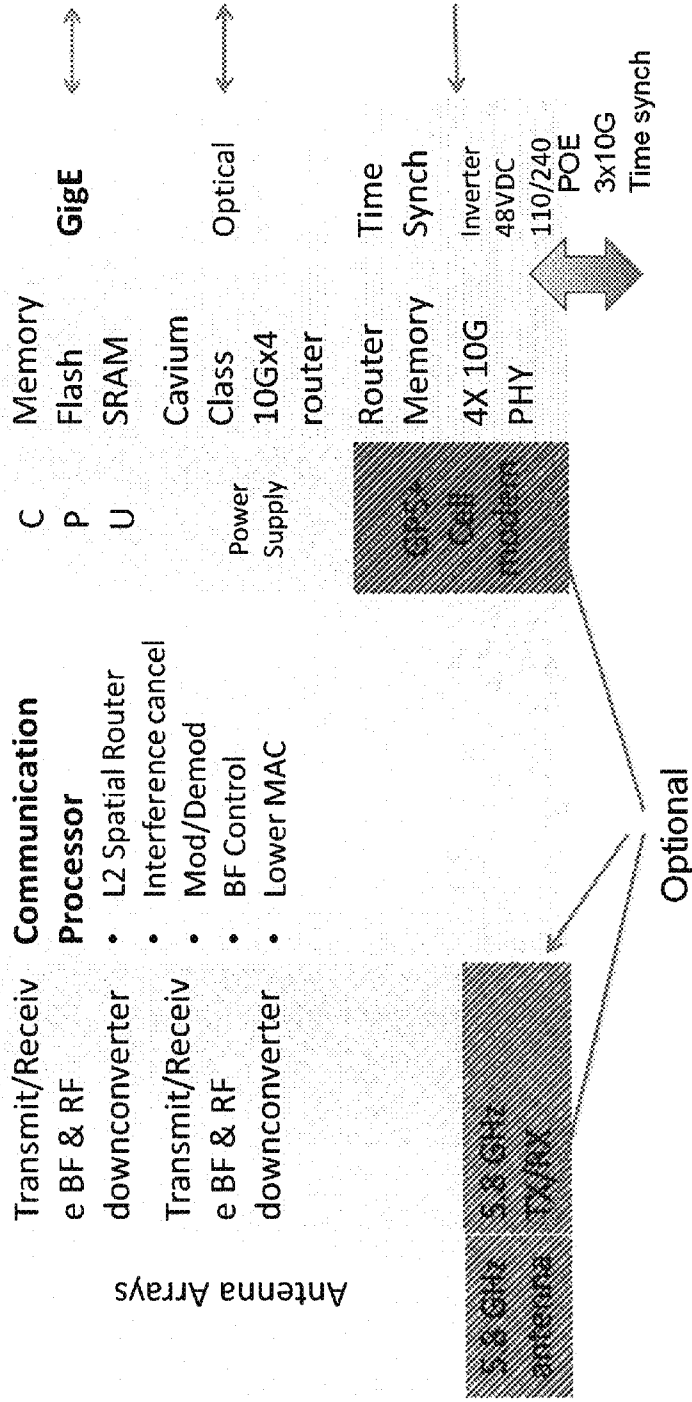


Figure 13: Antenna



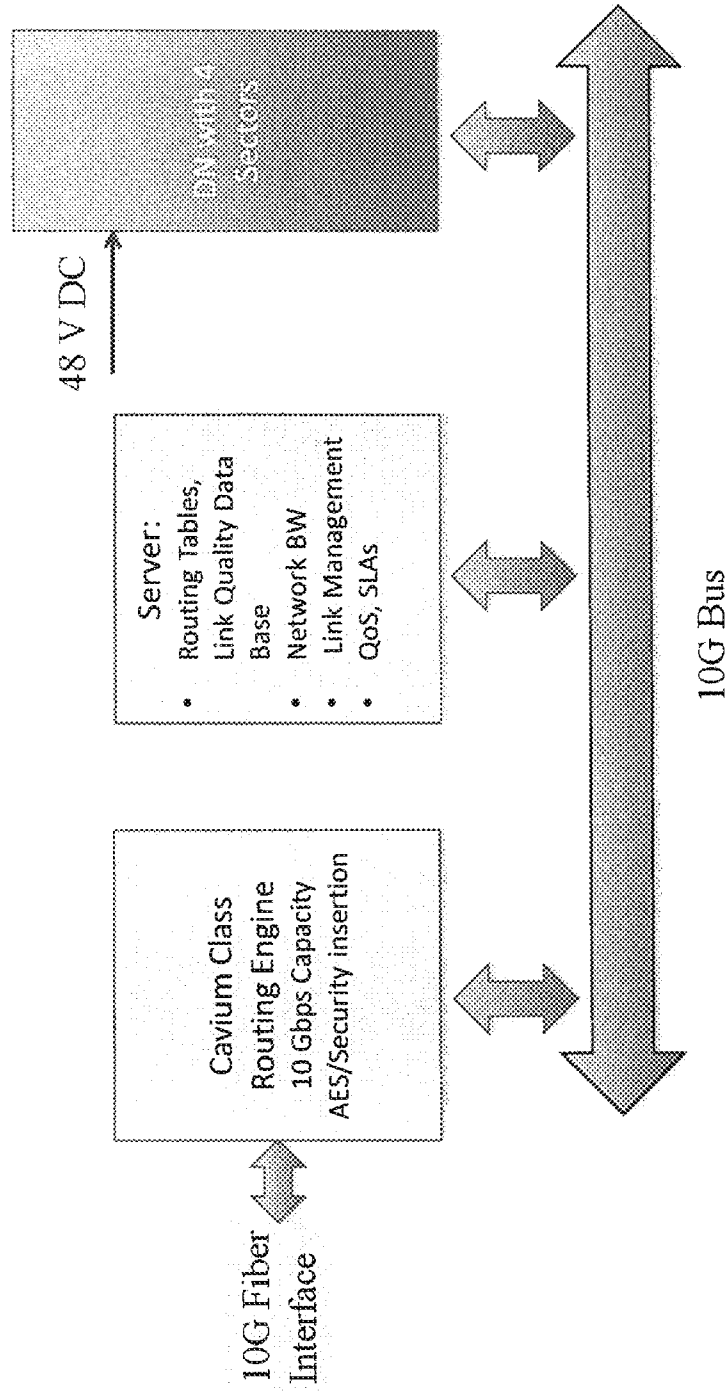
2 antenna arrays per sector
Antenna arrays can be 1d as shown
or 2 d
Element spacing is approx $\frac{1}{2}$ wavelength,
For example, can be $\frac{1}{3}$ wavelength also

Figure 14: DN Primary Sector



Other sectors don't have optical I/O & cavium (routing/switching) processor
 Primary sector is the master for packet switching and timing synchronisation and control

Figure 15: DN + Fiber I/O



Inovi proprietary & confidential

Figure 16: RF Maps Characterize Reflected Paths for Transmit & Corresponding Receive Beam

- Beam Coefficients for Receive and Transmit
- Parameter variance for tuning
- Signal Modulation
- Error correction
- Link Signal Loss
- Transmit Power Level

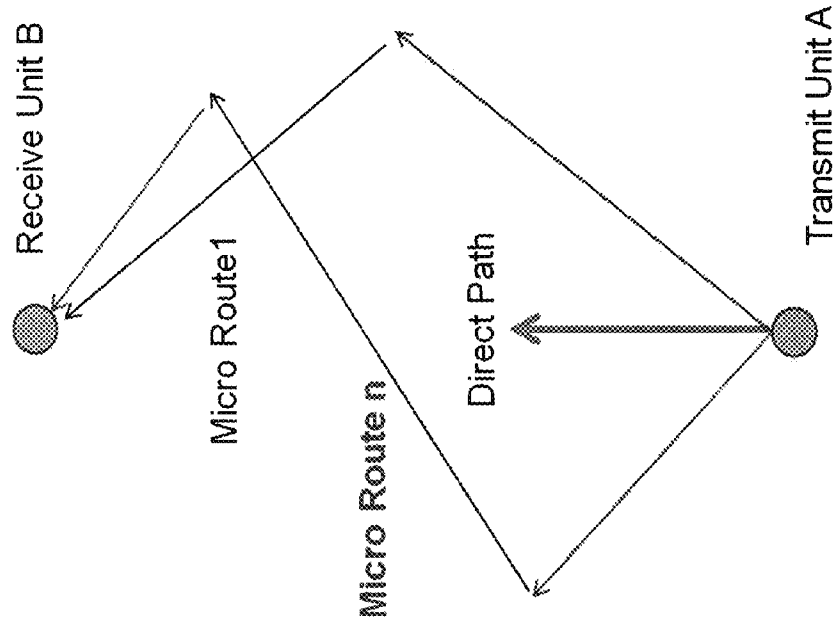


Figure 17: Data Collection for RF Maps

Blues receive, for each angle on red, blues cycle thru 360 degrees recording power received for each angle data sent to cloud

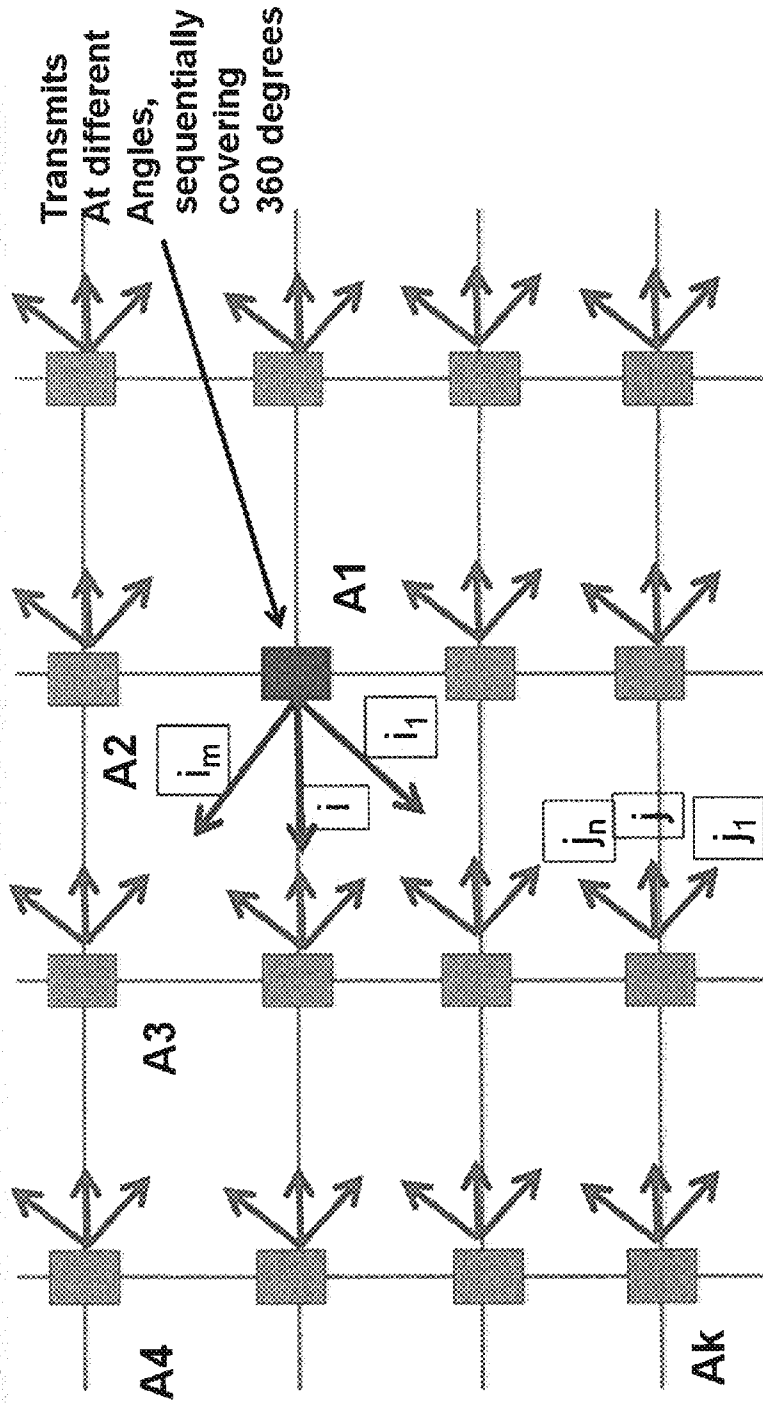


Figure 18: Predict Rain Micro Cells Motion to Reroute traffic around them

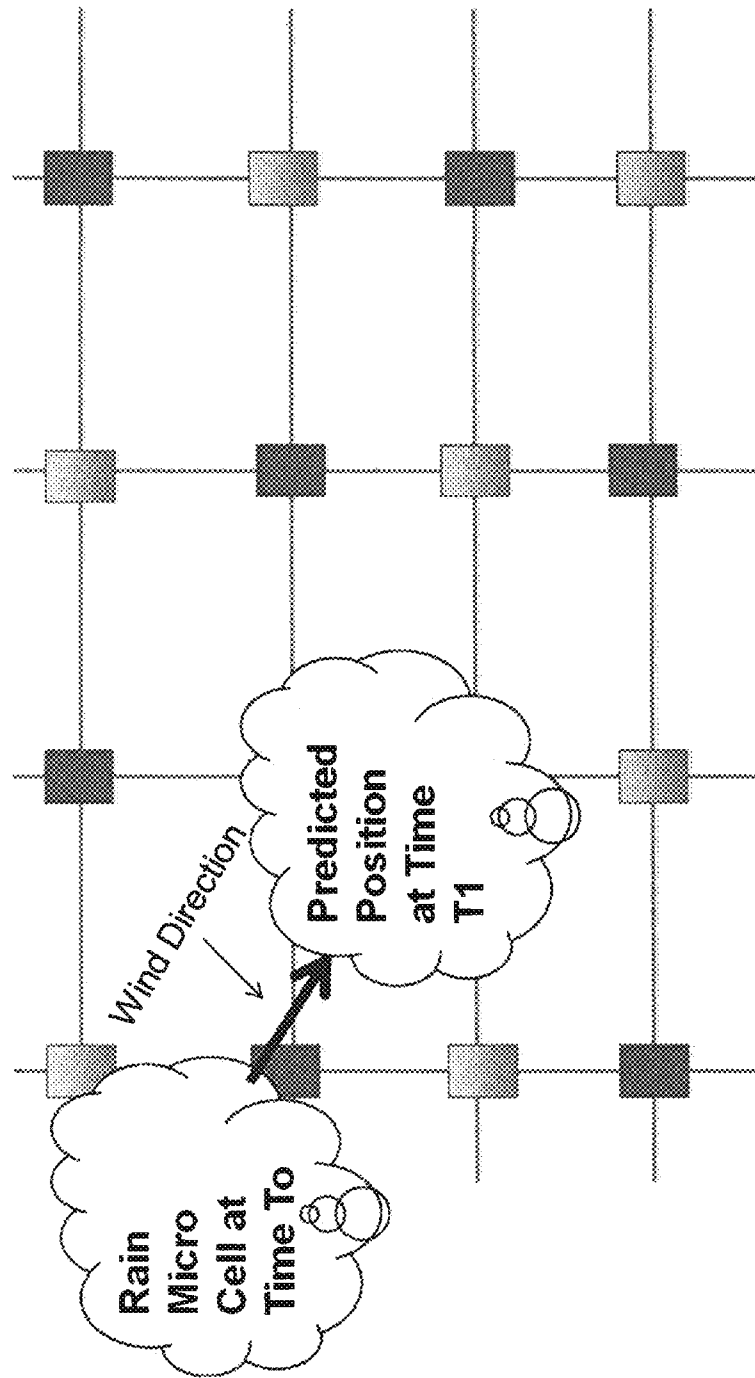
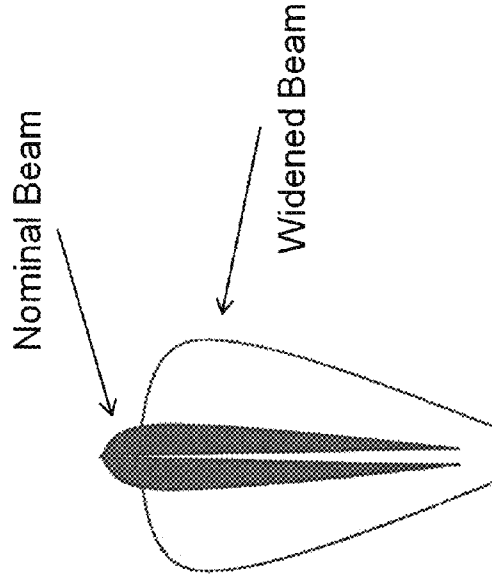


Figure 19: Beam widening



Parameters are

- Beam Coefficients
- Transmit Power
- Modulation
- Coding

CLOUD CONTROLLER FOR NEXT GENERATION DATA NETWORK

CROSS-REFERENCE TO RELATED APPLICATION

The present application is related to and claims priority under 35 U.S.C. 119(e) to U.S. provisional patent application Ser. No. 61/789,188, entitled "Next Generation Data Network," filed on Mar. 15, 2013. The contents of the provisional application are hereby incorporated by reference in its entirety.

TECHNICAL FIELD

The present disclosure relates generally to wireless communication. In particular, the present disclosure relates to methods and systems for controlling a multi-hop mesh wireless data network.

BACKGROUND

Data networks have traditionally been wired, either copper or fiber. Wireless spectrum has typically been used for subscriber links, not Data Infrastructure. The advent of micro cells, pico cells, nano cells, and hotspots has made the provisioning of wired Data prohibitively expensive. This is a result of inflexibility of wired networks; cable has to be run to add new points of presence. In many cases, this requires that streets be torn up.

Wireless Data networks can be much more flexible, and hence more economically viable. However, existing wireless Data network architectures limit this flexibility in several ways. One common form of wireless Data networks uses point-to-point radio links, typically operating in the microwave or millimeter wave bands.

Another common form of wireless Data networks use a cellular architecture. These networks are based on frequency reuse. Typically a cellular reuse scheme where the total available frequency spectrum is divided into sub bands, "colors", which are reused in disjoint geographical areas, "cells". For example, reuse factors of 4, 7, and 12 are popular. Other systems reuse time, polarization, or code. The disadvantage of all of these fixed, "cellular", reuse schemes are that they are inefficient, resources are wasted. For example, in a 7 to 1 frequency cellular reuse system, only 1/7-th of the total available frequency spectrum is available in a single cell, even if none of the adjacent cells are being used. Another disadvantage of these frequency reuse schemes is that it is difficult to add more cells in an area of high demand.

SUMMARY OF THE INVENTION

This invention is a wireless Data network that reuses the same frequency band in all cells. For example, if the frequency band is divided into three frequency channels, all three of these channels are available to every node in the network. Frequency channels are dynamically assigned by the Cloud to maximize network capacity taking into account potential interference from both other links in the network and from other sources. This is distinct from previous systems which try to maximize $C/(N+1)$, ratio of carrier power to combined noise plus interference power. Further, the system dynamically adjusts route and frequency channel assignments, transmit power, modulation, coding, and symbol rate

to maximize network capacity and probability of packet delivery, rather than trying to maximize the capacity of any one link.

The Next Generation Network is a time domain duplex (TDD) system (it can also be Frequency domain duplex, however the TDD is the more complex case, hence explained in detail). All sectors of each node transmit or receive at the same time. They are synchronized by a master controller. This allows the use of a single frequency across all sectors and mitigates adjacent channel interference. It is a multipoint system where there is no limit to how many nodes a node can communicate with. Each sector on each node has a 2D or 3D beamforming capability. This permits the unit to form a high gain beam that is pointed at the destination node and packet sent. The receiving node does the same it forms a high gain narrow beam pointed at the transmitting node. Beam coefficients are grossly computed in the cloud and fine tuned between the units.

The network uses two forms of route adjustment. A mesh network topology is used to allow traffic to be routed between two nodes via at least two geographically independent paths. This provides load balancing, so that when one path becomes congested or degraded, traffic can be routed around it. It also mitigates the effects of external interference. The second form of route adjustment is "micro-routes". Rather than just pointing the receive antenna at the transmitter, or pointing the transmit antenna at the receiver, this invention takes advantage of multipath to use indirect (non-line-of-sight) paths in addition to the direct (line-of-sight) path. The ability to use these paths provides several advantages. One is interference avoidance. Any interference that was arranging along the main beam will be mitigated by pointing away from it. Another is obstacle avoidance. The micro route can be used to mitigate obstacles, such as trees, signs, trucks, or similar that blocks the direct line-of-sight path. Further, by pre-computing these reflected routes they can be applied in real time.

Frequency channel assignments, transmit power, modulation, coding, and symbol rate for each link, both direct path and micro route, are dynamically optimized in the Cloud to maximize network capacity. Each node transmits probe packets to each of its neighbor nodes that can be n layers deep using a synchronous scheme. During a time interval known to its neighbors, the node transmits probe packets in various directions with a known pattern. The neighbor nodes listen for each of these transmissions in various directions, and record the received channel measurements. These measurements are periodically uploaded to the Cloud and processed. The processing identifies the viable micro routes and determines the range of link parameters available for each. This processing takes into account the received signal strength of the probe packets in each receive direction for each transmit direction. It also takes account of the reverse measurements, those made by the other node with its receive direction pointed in the previous transmit direction and with first nodes transmit direction pointed in the direction of the previous receive direction. Data is processed over time to identify static routes, those that are available over long periods of time, and periodic routes, those that are repeatedly available at certain times of day.

The received signal strength data collected by each unit is sent to the Cloud server for processing. This data consists of transmit angles, which in one embodiment of the invention are represented by beamforming parameters, and receive angles, which in the same embodiment of the invention are represented by beamforming parameters, and received signal transmit power.

A data processing function in the Cloud processes the data to identify micro-routes, reflected paths and interference levels between units. In an alternate embodiment of the invention, the processing is performed in each unit. In another alternate embodiment, the processing is performed in selected units. For each micro-route, the Cloud processor determines a set of signal parameters to be used for that route. These parameters are the transmit power level, beam pointing/forming coefficients plus modulation and coding. In one embodiment of this invention, transmit constellation pre-distortion parameters are included in the signal parameter set.

The RF mapping function is performed every 60 seconds. In alternate embodiments of this invention, it is performed at different rates. In one alternate embodiment, it is performed whenever capacity utilization is low. The RF mapping data is collected incrementally, distributed over time in different sections of the network. In another embodiment of the invention, the two units at the end points of a micro-route negotiate evaluations of beam pointing perturbations to optimize performance over time.

Meteorological data (weather forecasts) of rain and wind storms is used to identify phenomena, which might impact RF link performance. When a link degrades due to rain, the meteorological data is used to predict the motion of the rain microcell. These predictions are used to route traffic around the rain microcell to accommodate the reduced link capacity.

In high wind conditions, or when the structural mount is flexible, the beams are widened to maintain lock when the antennas are moved by the wind. To compensate for the widened, lower gain beams, the Cloud also reduces the data rate on these links. Wind direction is used by the Cloud to anticipate the need for beam widening.

Coarse beam parameters are computed on the server. These are geometric in nature and downloaded to the unit. Further the server can also compute perturbations on these parameters that the unit can try and then categorize by performance.

The Next Generation Network mitigates external interference in several ways:

1. By changing frequency channel
2. By using a reroute (route diversity) to avoid an interfered link
3. By using a micro-route to “look away” from the interference
4. By using a spatial interference canceller
5. By using a spatial interference canceller and a micro-route if the interference is along the original path or if additional interference suppression is required
6. By using adaptive modulation and coding
7. By using adaptive modulation and coding with any of the other techniques

BRIEF DESCRIPTIONS OF THE DRAWINGS

FIG. 1 shows a topology of the next generation data network, in accordance with one embodiment of the present invention.

FIG. 2 shows the deployment of the distribution nodes (DNs), client nodes (CNs), and a cloud controller (CC) in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 3 shows the four sectors of a distribution node, in accordance with one embodiment of the present invention.

FIG. 4 shows the next generation data network split into subnets each associated with a subset of fiber connected nodes, in accordance with one embodiment of the present invention.

FIG. 5 shows a transmit/receive exception pair characterized by the RF map function as an undesirable path in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 6 shows a micro route around a physical or interference blockage in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 7 shows a re-route around a physical or interference blockage in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 8 shows interference suppression using a micro route in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 9 shows the beamforming characteristics of a spatial interference canceller used in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 10 shows packet switched beams transmitted by a distribution node of the next generation data network, in accordance with one embodiment of the present invention.

FIG. 11 shows the block diagram of the four sectors of a distribution node of the next generation data network, in accordance with one embodiment of the present invention.

FIG. 12 shows using beamforming to maximize the coverage region of the side lobes for communication in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 13 shows the elements of the antenna arrays of a sector of a distribution node or a client node in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 14 shows the block diagram of a sector of a distribution node of the next generation data network, in accordance with one embodiment of the present invention.

FIG. 15 shows a bus interconnection between the sectors and fiber I/O in a distribution node of the next generation data network, in accordance with one embodiment of the present invention.

FIG. 16 shows characterization of the RF environment for reflection using RF maps in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 17 shows data collection of transmit angles, receive angles, and received power from the transmitter and receivers of a next generation data network for generating RF maps, in accordance with one embodiment of the present invention.

FIG. 18 shows prediction of rain microcell to reroute traffic around the rain microcell in the next generation data network, in accordance with one embodiment of the present invention.

FIG. 19 shows beam widening to maintain lock on a particular link of the next generation data network, in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION

The Next Generation Data Networks topology is shown in FIG. 1. The network (FIG. 2) consists of Distribution Nodes (DNs), Client Nodes (CNs), and a Cloud Controller (CC). The Distribution Nodes (FIG. 3) may connect to a fiber point of presence to get or send data from the “cloud”. The distribution nodes can talk with other distribution nodes and client nodes. The DN has one to four sectors. The Client Nodes have one or two sectors. The Cloud Controller manages routing and network performance analysis. A network may have multiple fiber connected nodes that can be used for redundancy and/or to distribute traffic entry/exit points. These networks can be further split into subnets each associated with a subset of fiber connected nodes (FIG. 4). These subnets are time

synchronized to minimize interference and edge effects between the subnets. Frequency channels may also be managed between these subnets to minimize interference between the subnets.

The Cloud Controller computes all of the routes from each Distribution Node to each of the Client Nodes and generates routing tables for each Distribution Node to/from each of its Client Nodes. These routing tables are downloaded into the nodes. The tables tell each node how to route a packet to its destination. Each node has its own MAC address. The cloud has a bird's eye view of the network as it collects statistics on each link to characterize it. These statistics are (but not restricted to) packet loss, obstructions and signal attenuation, structure rigidity, available of single or multiple micro routes. Using this data it computes predicted probability of packet delivery for each MAC address serviced by that DN. Another option is for the CC function to download all routes to the fiber connected DNs, and to each DN-incremental routes to destination MAC addresses and micro routes relevant to that DN. An example of the schedule is:

Macro Routing (Once/10 Sec)

Cloud computes all possible routes between fiber nodes and MAC addresses;

Uplink & downlink;

Cloud Ranks these routes;

Routes are weighted with link quality, speed, time of day, and traffic congestion;

For each MAC address we have primary and several alternate routes that are pre computed with a cost function;

Cloud refreshes ranking—no faster than once every 10 seconds;

Refresh triggered by a substantial change in link quality;

Refresh includes transmit parameters of links;

Route probe commands;

The cloud keeps alternate routes alive with pings and pre-scheduled slots.

Micro Routing (Once/Hr)

Cloud refreshes micro routes & ranking every hour based on RF map updates & Route Probes.

Exception Pairs (Once/Hr)

Exception transmit receive pairs for SFNs are recomputed every hour.

The other functions performed by the cloud controller is creation of RF Maps, management and command of data collection by nodes, analyses of link performance data collected by the nodes, determination of transmit power levels for the nodes and clients that are dependent on the destination of the packet. The cloud controller also performs management of the single frequency network, virtual networks, sub nets, and coexistence with other networks. The cloud controller may also manage interference at the node level or have the fiber connected node and the node itself manage these functions depending on timeliness of a response from the cloud controller in response to changing conditions.

When a fiber connected node receives a packet, either from the fiber for a Distribution node, or from a client for any node, it finds the destination MAC address in its routing table, and appends the routing header specified in the table to that packet. The packet is then transmitted to the next node in the header. The header entry tells the node which sector and beam number in that sector to use. Prior to transmission the sending node strips the header entry it just used from the packet to expose the next entry which will be used by the receiving node.

To implement load balancing, the Cloud Controller includes two, or more, entries in the routing table for a given destination node. Each entry has an assigned probability. The sum of the probabilities for a given destination is always one.

Thus if there is only one entry, the assigned probability value is 1. If there are more than two entries for a given destination, the sending node generates a random value between 0 and 1. If the value is less than the probability assigned to the first entry, that entry is used. If not, the value is compared to the sum of the probabilities assigned to the first two entries. If the value is less than this sum, the second entry is used. If there are more than two entries, this procedure is repeated until the packet is sent. Further, the rules or cost function used for routing can consist of (but not restricted to)

No Exception Routing Pairs (SFN)

For SFN operation we create transmit and receive exception pairs (FIG. 5) derived by the RF Maps function. That is, if a transmitter is sending a packet to a client/DN with a specific beam type (main or micro route), that may interfere with other units. We ensure if these units are receiving at that time, that their main beams are not pointed in the interference direction. This can be achieved in two ways—one is leave the receive slot empty or option 2 (the preferred option) is to order the packets so that the receiver and transmitter have their beams occupied in different directions. Exception pair may also be created to prevent overlapped transmission beams of two adjacent sectors of the DN from interfering with one another.

Only Terminal Traffic on Low Reliability Links

On low reliability links where the net data throughput is low we may want to restrict traffic over those links to traffic where the destination MAC addresses are only reachable by going over those links. Further, we leverage the fact that traffic density on a link decreases exponentially on average as we move away from the fiber connected node. This gives us excess unused capacity on most links that are 1 or more hops away from the fiber connected node.

Lower Modulation/Higher Error Coding on Low Reliability Links

We can further enhance the performance of these links by slowing down the payload rate to ensure timely packet delivery by reducing the modulation complexity (qam 16 to qpsk for example) for example and/or increase the error coding (Rate $\frac{3}{4}$ to rate $\frac{1}{2}$ code for example).

Lower Modulation on Sparse Data Traffic Links

On links with sparse data, that is a lot of unused capacity we can reduce the modulation and increase the error coding to achieve a higher reliability of packet delivery.

Double Packet Routes of Trunk Data Over Lossy Links

High priority packets and "trunk" data (that is data that needs to be forwarded on) can be sent over multiple routes simultaneously. Further large packets can be broken down into smaller packets as that may preserve the integrity of the packets that are being received from different directions.

Maximization of Probability of Packet Delivery

The above mentioned technique may be used in addition to algorithm choosing links for the route that maximize the probability of packet delivery. The fiber connected node and the cloud routing & processing function are compiling statistics on each link that are being used to predict future performance of each link.

A Minimization of Hops

If the probability of packet delivery is the same on two routes, and there is no congestion, the route with the lower number of hops will be chosen.

The Next Generation Data Network implements two forms of route adjustment. The network topology allows packets to

be routed between two nodes via at least two spatially independent paths. This provides load balancing, so that when one path becomes congested or degraded, traffic can be routed around it. It also mitigates the effects of external interference. The second form of route adjustment is “micro-routes”. Rather than pointing the receive antenna at the transmitter, or pointing the transmit antenna at the receiver, this invention takes advantage of multipath to use indirect (non-line-of-sight) paths. The ability to use these paths provides several advantages. One is interference avoidance. Any interference that was aligned with the main beam is mitigated by pointing away from it. Another is obstacle avoidance. The micro-route can be used to mitigate obstacles, such as trees, signs, trucks, or similar objects that block the direct line-of-sight path.

Frequency channel assignments, beam forming coefficients, transmit power, modulation, coding, and symbol rate for each link, both direct path and micro-route, are dynamically optimized by the Cloud Controller to maximize network capacity. The two nodes have the option of further optimizing the cloud computed beam parameters to include multipath effects, unit misalignment, calibration & temperature dependent errors, etc.

Each node transmits probe packets to each of its four neighbor nodes using a synchronous scheme. During a time interval known to its neighbors, the node transmits probe packets in various directions with a known pattern. The neighbor nodes listen for each of these transmissions in various directions, and record the received channel measurements. These measurements are periodically uploaded to the Cloud Processor and processed. The processing identifies the viable micro-routes and determines the range of link parameters available for each. This processing takes into account the received signal strength of the probe packets in each receive direction for each transmit direction. It also takes account of the reverse measurements, those made by the first node with its receive direction pointed in the previous transmit direction and with second node’s transmit direction pointed in the direction of the previous receive direction. Data is processed over time to identify static routes, those that are available over long periods of time, and periodic routes, those that are repeatedly available at certain times of day.

To mitigate the impact of lossy links, retransmissions are used. The receiving node is capable of receiving fragments of a packet over time on the same link or distributed over other links and reassembling them. When lossy links are detected, the sending node fragments the packet and sends copies of fragments using different routes to achieve route diversity and increase the probability of packet delivery. The node can tradeoff the losses between all of the entries in the routing table for a given destination node, including the direct-route and a micro-route. If the losses on the micro-route are lower it may choose to use the micro-route for packet delivery.

For a link that is unresponsive, due to a physical blockage or radio interference, on the path, the transmit node tries to establish communication on a “micro-route”. It may try multiple micro-routes. If that fails it returns the packet (or packet number) to the last unit (DN_{n-1} unit) that sent it.

For rerouting, FIGS. 6 and 7, the DN_{n-1} node has a list of alternate routes for each destination node pre-computed by the Cloud Controller. The node modifies the routing header using one of these and retransmits the packet. It is entirely possible that the packet (or packet identifier) may retrace several steps before it gets on a new route.

Further, this unit deletes routes that contain the failed route from its list. This is also done at the fiber connected node and all other nodes in the network. Only the cloud computation

can reestablish this route. The deletion can be done based on persistence or on just one occurrence of failure.

All communication failures, link SNR etc. are communicated back to the fiber node and the cloud.

Each sector of the node is equipped with an FFT, which runs over the entire spectrum creating a power spectral density map or a “spectrum analyzer”. This data is used by the Cloud Controller to select the frequency of the link. The transmit and receive frequencies can be different and are determined by the interference at the receiver rather than at the transmitter. Further each sector of the node can operate on a different frequency.

For each fiber point or broadband point of presence (a high bandwidth link for taking and bringing data from carrier/internet) and each destination MAC address on the network, the Cloud pre-computes the possible routes with a cost function. The net of the cost function is characterized in data rate, delay, and reliability of packet delivery. It also takes into account the availability of multiple routes between two units using micro routes.

Once a unit determines it cannot communicate on its primary path it utilizes pre-stored micro routes or an alternate path computed by the Cloud. If these fail it sends the packet back to the DN_{n-1} . DN_{n-1} has a list of pre-computed routes for a destination mac address. DN_{n-1} replaces the packet routing header with one of the pre-computed routes that does not utilize a “failed” route.

The fiber connected node and the server are informed of the failed route or the use of micro-routes and any other route anomalies. They are also given information on links that are operational on signal level, signal headroom, fading margin etc.

The Cloud computes all routes from the Distribution Node to the destination node. These are downloaded into the fiber connected unit. In an alternate embodiment of the invention, they are downloaded into an auxiliary units attached to the fiber connected unit. When the fiber connected unit receives a packet for a MAC address it adds the routing header on it, and sends it to the appropriate adjacent unit. The routing header tells a unit as to which sector to transmit the packet, and the beam “number” on that sector. The unit strips the parameters from the route header exposing the next one for the receiving unit. The fiber connected unit will round robin through equal cost or near equal cost routes for load balance on the network.

If a communications failure occurs, the node first retransmits the packet. If that fails, it tries a micro-route or an alternate route computed by the Cloud. If that fails, it returns the packet to the last unit (DN_{n-1} unit) that sent it. The DN_{n-1} unit has a list of alternate routes for each destination node, pre-computed by the Cloud Controller and downloaded to that unit. It modifies the routing header using one of these and retransmits the packet. This process is repeated until the packet is either returned to the Distribution Node, or it is delivered. If it is returned to the Distribution Node, that node re-sends it on a different path.

When a failure occurs, the sending unit may delete the failed primary path, the alternate path, and/or the micro-route from its routing table, and propagates this information to the other nodes in the network, including the fiber connected node. The sending unit may also restrict the use of the primary path, certain alternate paths, and/or the micro-route between the updating of the routing table by the Cloud. The fiber connected node reports this information to the Cloud. Only the Cloud processing can reestablish this route. In an alternate embodiment, instead of deleting a micro-route based on a

single failure, a failure counter is maintained for each micro-route, and only when a predefined threshold is exceeded, is the micro-route deleted.

The means of detecting a static reflector are to collect reflection data at different times of day or days of week or month or year. If the reflected path appears in all of them accounting for the variance of beam pointing errors, it is deemed to be a static reflector. Dynamic reflectors that have a known periodicity may also be used but as a function of that periodicity.

By using a micro-route, rather than the direct path, the Next Generation Network changes the arrival angle of the signal at the receiver. This allows the network to maximize the signal to interference ratio at the receiver. For example, typical side-lobe to main-lobe gain is -13 dB, and a typical micro-route loss is 3 dB. In this case, the improvement in C/I is 10 dB. (FIG. 8)

Micro Route+ Modulation

In this case we use the microroute and reduce the modulation & effective data rate on the link. For example if we increased the margin by 6 dB from modulation-data rate then our effective margin over the interferer is 16 dB.

Micro Route+ Interference Cancellation

We use a spatial interference canceller (FIG. 9). The canceller does an “inverse beam former” or creates a null in a direction. We have shown that these nulls maybe as deep at 28 dB over and above antenna isolation. In this case we have an effective margin of 38 dB over the interference.

Micro Route+Interference Cancel+Modulation

In this case if we had a 6 dB advantage from modulation-data rate our effective margin is now 44 dB.

Single Frequency Networks

We enable the complete network to be run on a single frequency channel rather than alternating frequencies. This enables the most efficient usage of the frequency spectrum. We achieve this via RF Maps. Using the RF Maps we:

1. Compute Power levels for the unit to transmit a packet (knowing its destination). Computed power levels reflect the real world power that is required to close the link at a desired reliability/availability level. It is important not to transmit at max power where we can use lower power to minimize self interference in the network.
2. Beam pointing coefficients coordinated with the packet transmission sequence across exception pair links. This eliminates the “exception pairs” generated by the RF Maps function
3. Synchronization and synchronization offsets for irregular deployments. In a regular or symmetric deployment alternate nodes transmit and receive. In an irregular deployment there may be multiple DNs on a path corresponding to one link on another path that need to be synchronized.
4. Managing C/I (Carrier to Interference ratio) at nodes via transmit power, modulation, packet sequence and beam pointing coefficients. The C/I is managed so that the packet can be delivered with a reliability level that has been guaranteed to the user.
5. Use 3D RF maps to compute off axis self interference combinations
 - a. Auto measure sidelobe performance;
 - b. Utilize the self interference cross correlation matrix to eliminate transmit receive combinations (beam weights);
 - c. Achieved via packet scheduling.

Virtual Networks

Multiple virtual networks can be created over the same hardware links. These networks can each be assigned data

bandwidth at the fiber node(s). The reason is that the fiber node determines total throughput of the network. Further, we can overbook each network if required using statistics that are available historically and that we have collected on the nature of the burstiness of the data. This will allow us to allocate bandwidth dependent on the guaranteed level of service for each virtual Network.

Network Coexistence

Multiple Operators may plan to use our equipment in the same area, and maybe on the same poles. Using communication between our cloud controllers we can manage the frequency, the synchronization of transmit and receive slots to minimize interference and maximize throughput. Further we can also coordinate packet sequencing between these networks to maximize antenna isolation both from beamforming and polarization. The SFN is key to making these networks operate as it increases the degrees of freedom that we have to operate both in frequency and/or polarization.

Network and SubNet Synchronization & Security

The network uses two levels of synchronization. Coarse synch using GPS to a nominal of 100 nanoseconds and a fine synch that using signaling between units. Sub nets are synchronized in the cloud and their transmit, receive, and beam pointing along the edges of subnets is coordinated and synchronized. Further subnet coordination also utilizes the sparse data or low data densities of the links at the periphery of each subnet to maximize data traffic in the core of each subnet. The core is defined as the fiber DN and the first hop DNs. Security leverages on the delays between units and time of travel between units. Using this as a key we “open” the packets only at the destination. Further due to the assignment of different routes along which the packet may travel, assigned at the fiber attached node the key is randomized. Further error bounds are placed to account for retransmits that maybe encountered along the way. Alternatively we can add “incremental” keys in each DN to account for retransmits.

Distribution Node (FIG. 3)

Describes nodes in a millimeter wave multi-hop network that provide high data rates (of the order of multiple Gigabits per second), flexible deployment (not requiring manual aiming of antennas) and robustness to multipath fading. FIG. 3 shows a node with 4 faces or sectors. All sectors transmit and receive at the same time. Each sector has a packet switched beam. That is the beamforming coefficients are applied to enable the sector to point the beam at the client or receiver of the packet to be transmitted. Further, the beam is then switched to point (FIG. 10) in the direction of where the next packet needs to be sent. Each sector has a table of MAC or IP addresses and beamforming coefficients that the sector is authorized to communicate with. There maybe multiple sets of beamforming coefficients for the main beam or a range assigned to each coefficient. Micro route beamforming coefficients are also included in this table. All 4 sectors are synchronized, one sector is the master and the other 3 sectors are the slaves (FIG. 11). All sectors transmit OFDM signals. However, the system can work with single carrier signals

The nodes can perform adaptive transmit and receive beamforming, and provide omnidirectional coverage as well as high antenna directivity. That is, each node can communicate with neighbors in any direction of its choosing in a 360 degree field of view, while being able to synthesize a highly directive beam in the direction of the node that it is communicating with at a given time. Link robustness is achieved by providing diversity against fading, which can occur due to multipath reflections from surfaces such as the ground, build-

ing walls, or vehicles. The throughput available at each node is increased by allowing for multiple simultaneous links with different neighbors.

The reason for multi-hop/mesh is that it permits us to “look” around corners/buildings and extend the distance we cover. By having multiple routes to the destination we also increase the effective availability/reliability of the network.

There are two node types, distribution nodes (DNs) and a client node (CNs). A distribution node has up to four sectors, each covering 6 degrees \times 90 degrees (we can change this coverage for 2D beamforming cases). The fundamental beam is approximately 6 degrees \times 3 degrees and is beam formed to sweep over the complete angle. Each sector has a beam that can be pointed \pm 45 degree. Further, we can reduce the coverage area but utilize the sidelobes to communicate with the close in clients who maybe on the extremities. For example (FIG. 12), we can reduce the actual antenna coverage to 6 degrees \times 45 degrees, and use the beamforming to maximize the coverage region (that is maximize the side lobes rather than main lobe). That is clients in the 22.5 degrees on either side of the coverage region are communicated to by the sidelobes. These links can be closed at a lower transmit EIPR as the clients are closer in and/or at a lower data rate. This enhances the link for the “trunk” connections. These nodes are reconfigurable for other angular coverage, which can include beam forming in elevation in addition to azimuth.

In an alternate embodiment, a two dimensional (2D) sweep is used. This 2D beam forming has an $n \times m$ degree beam that can be steered vertically in elevation and horizontally in azimuth. This embodiment is used for building top to building top as buildings are at different heights. It is also used to illuminate the side of a building from the street or top of another building.

The antenna array has 32 elements, each 6 degree by 90 degree (FIG. 13). In alternate embodiments, it may have 16, 12 or any other number of elements. There are two such arrays per sector. On transmit, both arrays are used with an Alamouti code that permits mitigating multipath effects and increases the transmit power via combining. On receive, both arrays are used with MRC combining.

A block diagram of the sector is shown in FIG. 14. It consists of the antenna arrays, which feed into the beamforming chips. These chips can be n element chips for simplicity we describe a 32 element beamformer. These beamformer chips use the coefficients for a packet on transmit or receive. We assume a single transmit/receive chip for each antenna array. We can have a separate transmit and receive chip with a switch or separate arrays for transmit and receive. These chips can output/input the combined RF signal at baseband or an IF frequency. This signal is modulated/demodulated by the communication processor which filters, combines signals from the two antennas for receive or generates two signals on transmit. It also has a QAM/QPSK mod/demod function with error coding. Other functions are shown in the FIG. 14 as an example. The MAC function is part of this processor—it can be in the baseband chip or a separate FPGA. Further Mac filtering takes place that is the routing header is looked at to determine where it goes and is sent to the appropriate sector directly or via the primary sector. We can use the 10 G bus as the interconnect bus or multiple 1 G buses or some other digital bus between the sectors. Connection to fiber I/O is shown in FIG. 15.

The multipath propagation characteristics for a highly directive mm wave link are quite sparse. Spatial diversity is provided by using multiple transmit and/or receive arrays offset in the horizontal and vertical directions. The horizontal offset provides diversity against wall reflections, while the vertical offset provides diversity against ground reflections.

The amount of offset is chosen based on typical link ranges and typical distances from scattering surfaces. Thus, each receive antenna array would generate a single mm wave signal. For diversity combining of signals from multiple receive arrays, these signals can be further combined by adjusting gain and phase (e.g., based on a maximal ratio combining strategy), or one of the signals can be selected (e.g., based on amplitude or power). This can then be sent to the single channel chipset.

Each node has a centralized controller that routes packets to and from the appropriate faces. Isolation between the radio frequency (RF) circuits for different faces is sufficient that all faces can simultaneously transmit, or simultaneously receive. By designing higher layer protocols appropriately, the need for a face to transmit while another face is receiving can be avoided, thus greatly simplifying the isolation requirements.

An additional means of increasing data rate and interference rejection could be to use polarization coding and reverse polarize the two antennas in the sector. We could take this approach for doubling the effective throughput I/O from the unit.

In one embodiment, output from the unit is GigE, SynchE; power into the unit is 48 V DC, 110 VAC, 240 VAC; and interconnects between the sectors is via a multiple gigabit Ethernet ports or a single 10 G port.

Other Functions and Components in the System

1. 3D RF Maps

We need to characterize the RF environment for reflections (FIG. 16), interference (FIG. 5), and link losses. Characterization is done using probe packets. While we discuss a two dimensional case, the maps are generalized for three directions. With vertical scanning (or elevation scanning on the beam) we can characterize the reflectors, interference and side lobes in three dimensions. For simplicity of understanding we discuss a two dimensional case.

To do this we start with a transmit beam on unit A and receive on all “surrounding” units. The number or depth of surrounding units is based on the propagation losses. Data collection is performed as follows (see FIG. 17).

Unit A1 transmits at angle i . All units receive at angle j . And then at $j+1, j+2, \dots n$;

Unit A1 then transmits at angle $i+1$, units receive at $j, j+1, \dots n$;

...

Unit A1 transmits at angle m , all units receive at $j, j+1, \dots n$;

Unit A2 transmits in the same sequence as A1 (above), all units receive the A2 signal at $j, j+1, \dots n$;

...

Unit Ak transmits at angle m , all units receive the Ak signal at $j, j+1, \dots n$.

The data collected on transmit receive are transmitted to the server in the cloud. The data consists of transmit angle or beamforming parameters and receive angle or beamforming parameters, received signal transmit power, multipath parameters if any.

Note transmit and receive signals are shown as “geometric” for simplicity of explanation; however these beams may be “amoeba shaped”. For example, amoeba shaped beams may be constructed from a combination of geometric beams from the multiple antenna elements (e.g., 32) of the antenna array. The amoeba shaped beams may be used to increase the probability of packet reception or to increase the data rate via independent paths. In one embodiment, amoeba shaped beams may be created between two nodes independent of the Cloud.

The data analytics function in the server or super unit or possibly in each unit, sorts through the data to compute reflected paths and interference between units. For each of these a signal power level is associated, a modulation parametric set is associated which may include distortion of the QAM constellation etc.

This data collection exercise can be performed on a periodic basis or ad hoc. The data collected can be incremental, that is we do not have to collect all data in one contiguous set but can be distributed over time. Further we will re-explore the reflected paths periodically, further the two units using the reflected path may further optimize beam parameters between them to optimize communication.

Whenever there is no traffic or periodically the units are set to scan one transmitting and one or several receiving. We correlate the transmit power and “angle” with what receivers see again, received power and angle. The angles can be in azimuth and elevation.

To make a single frequency network we use the RF Maps to pin point interference between cells. As both “cells” have beam forming we create “packet exclusions,” that is if a packet is being sent to a client then the interfered node will not look (receive/transmit) in that direction while this operation is in progress.

RF Roads are angular Transmit—receive combinations.

Discover RF roads between units using the scanning beams;

RF Roads have data speed, reliability, and exclusion zones; For desirable roads—permanent reflectors are sourced by collecting data repeatedly over time.

2. Weather & Environment Tool

Using Met data we are aware of rain & wind storms that may degrade the performance of our links. When we see a link degrade due to rain we can predict the motion of the rain microcell via met data (wind direction & speed). This prediction is used to route traffic around the rain microcell and link capacities in the rain micro cell will be reduced (FIG. 18).

For wind we collect data on degradation of the beam—that is we may have to widen a beam to maintain lock on a particular link. We can use this data that we have collected historically, to predict effects in a wind storm which will require that we predict and widen the beam (FIG. 19), and reduce the data rate to maintain the link budget. Further we can use wind direction in the storm to predict specifics of beam parameters and how they are affected. (Beam may only need to be widened for particular wind directions).

3. Beam Parameter Computation

Coarse beam parameters are computed on the server. These are geometric in nature and downloaded to the unit. Further the server can also compute perturbations on these parameters that the unit can try and then categorize by performance. Initial beams may be computed using the known GPS locations of the two end of the link and approximate direction of the sector face. By using multiple measurements to various destinations that the sector can communicate with, the Cloud Controller can compute the “nominal” direction of the sector (that is at 0,0).

When a new unit is added to the network its geographic location is entered into the provisioning system. The appropriate DNs and their appropriate sectors (using GPS data) will look for the new unit. The new unit is in receive & scan mode unit it acquires the invitation to join the network from a DN. AES keys may be exchanged before it is permitted to enter the network. Further, it will receive the definition of transmit slots, etc. from the DNs.

The perturbations can be for enhanced performance for multipath (reflectors have been characterized in RF Maps) or

for interference cancellation, or both. The interference cancellation nulls can be pre-computed based on beam directions and cancellation nulls in other directions. Perturbations for antenna imperfections are also taken into account and calibrated in the Cloud Controller. This is similar to the process to calibrate misalignment or the mounting angle of the sector.

Transmit or receive beamforming can be performed based on any of a number of well-known mechanisms, including beamforming at radio frequency (RF) or intermediate frequency (IF) by controlling the relative gains and/or phases of each element, or beamforming in complex baseband by controlling the complex gains of the baseband signal to/from each element for transmit/receive beamforming.

The beamforming and diversity functionalities can be modularized and implemented in a number of ways. As one example, sophisticated integrated circuits developed for short range (e.g., indoor) applications can be leveraged for the longer range outdoor applications considered in this disclosure. As FIG. 14 show, a single channel chipset that handles up/down conversion, baseband processing, and link layer/medium access protocols, can be used for each face. Alternatively, Beamforming chipsets customized for the outdoor application can be layered on top of single channel RF up/downconverter chips. In order to control line losses, it is advantageous to place a beamforming integrated circuit (IC) close to the antenna array that it is controlling. For transmission, a transmit beamforming IC would take the mm wave signal from the single channel chipset, replicate the signal for each transmit element, provide power amplification, and modulate the gain and phase to beamform in the appropriate direction. The amount of power amplification would be set so as to adhere to regulatory guidelines on radiated power, while overcoming line and coupling losses from the single channel chipset to the antenna elements. For reception, a receive beamforming IC would provide low noise amplification as close to the antenna elements as possible (so as to improve the overall noise figure), while adjusting the gain and/or phase of the received signal at each element before summing these signals to implement receive beamforming.

4. Routing Tool

For each fiber point or broadband point of presence (a high bandwidth link for taking and bringing data from carrier/internet) and each destination MAC address on our network, we pre-compute the possible routes with a cost function. The net of the cost function is characterized in data rate, delay, and reliability of packet delivery. It also takes into account the availability of multiple routes between two units using micro routes.

Further, for rerouting/routing failures, once a unit determines it cannot communicate on its primary path it utilizes pre-stored micro routes or an alternate path from the routing table. If these fail it sends the packet back to the DN_{n-1} . DN_{n-1} has a list of pre-computed routes for a destination mac address. DN_{n-1} replaces the packet routing header with one of the pre-computed routes that does not utilize a “failed” route.

The fiber connected node and the server are informed of the failed route or the use of microroutes and any other route anomalies. They are also given information on links that are operational on signal level, signal headroom, fading margin etc.

Routing is computed based on RF parameters/statistics of each link and traffic on the link, to maximize probability of packet delivery at its destination. The routing tool can also order increased time coding of packets for error recovery, usually applied to high priority/critical packets and also for lossy links. For destinations that can only be reached via excessively lossy links the routing tool may order two or more

15

copies of the packet to be sent on distinct routes, where the receiving client has the responsibility to reconstruct the original packet from error corrected fragments of the original packet. Further for excessively lossy links the routing tool may order a large packet be broken down into smaller sub packets.

Routing header may include:

Route number, sector number, beam number, power to transmit at, modulation and error correction parameters, number of bits;

Route number+1, sector number, . . .

. . .

Ethernet port address;

Data;

. . .

Data.

5. Micro Route;

Micro-route is a reflected route, pre-computed in the cloud using "RF Maps" The parameters (beam angle, modulation, coding, tx power, . . .) for these routes are downloaded to the unit and then fine tuned by the unit. Micro routes utilize "static reflectors" who properties are known as a time of day or are static over all time.

The means of detecting a static reflector are to collect reflection data at different times of day or days of week of month or year. If the reflected path appears in all of them accounting for the variance of beam pointing errors, it is deemed to be a static reflector. Dynamic reflectors that have a known periodicity may also be used but as a function of that periodicity.

Micro route information also includes beam parameters, data rate, modulation, error correction and other parameters to let the units know how to close the link at a required reliability level. Multiple micro routes between two units can be downloaded and prioritized.

The other usage for micro routes is for interference management. By changing the angle of arrival of the signal it can present a side lobe of the antenna gain pattern to the interference.

Micro routes are periodically tested along with all other routes by the Cloud Controller.

I claim:

1. A method for controlling a multi-hop mesh wireless data network by a controller, comprising:

characterizing by the controller data on radio frequency (RF) links at a plurality of nodes of the wireless data network;

analyzing the data on the RF links to compute parameters for one or more alternate routing paths from a source node to a destination node in one or more hops across the wireless data network;

loading the parameters for the alternate routing paths to the source node to direct the source node on how to route a data packet from the source node to the destination node through any intermediate nodes;

receiving statistics on the delivery of a plurality of data packets from the source node to the destination node; and

adaptively adjusting the parameters for the alternate routing paths in response to the data on the RF links and the statistics on the delivery of data packets.

2. The method of claim 1, wherein said characterizing comprises controlling a pair of the plurality of nodes to collect signal strength data on the RF link between the pair of nodes.

16

3. The method of claim 2, wherein the pair of the nodes comprises a transmitting node and a listening node, wherein the listening node is located one or more hops from the transmitting node.

4. The method of claim 3, wherein said characterizing further comprises the controller commanding the transmitting node to transmit a probe packet in a plurality of transmit directions around the transmitting node, and commanding the listening node, for each of the transmitting directions of the probe packet, to collect received signal strength data of the probe packet in a plurality of incident directions around the listening node.

5. The method of claim 4, wherein the probe packet is transmitted and received using a beam having a geometric shape.

6. The method of claim 4, wherein the probe packet is transmitted and received using amoeba shaped beams, wherein the amoeba shaped beams are constructed from a combination of geometric shaped beams from a plurality of antenna elements.

7. The method of claim 1, wherein said analyzing the data on the RF links comprises detecting one or more multi-path reflectors that reflect an RF signal of an RF link between a pair of the plurality of nodes.

8. The method of claim 7, wherein said analyzing further comprises detecting the multi-path reflectors over a period of time and categorizing the multi-path reflectors as static.

9. The method of claim 7, wherein the computed parameters comprise beam parameters for forming a transmission beam and a receiving beam that take advantage of the multi-path reflectors in a non-line-of-sight path between the pair of nodes.

10. The method of claim 7, wherein the computed parameters comprise a transmit angle for a transmission beam and a receiving angle for a receiving beam that take advantage of the multi-path reflectors in a non-line-of-sight path between the pair of nodes.

11. The method of claim 7, wherein said analyzing the data on the RF links further comprises detecting an interference source that interferes with a RF link between a pair of the plurality of nodes.

12. The method of claim 11, wherein the computed parameters comprise a transmit angle for a transmission beam and a receiving angle for a receiving beam that take advantage of the multi-path reflector in a non-line-of-sight path between the pair of nodes to mitigate the interference source, wherein the receiving angle is offset from an angle of arrival of the interference source so as to increase a ratio of signal to interference.

13. The method of claim 1, wherein said analyzing the data on the RE links further comprises detecting an interference source that interferes with a RF link between a pair of the plurality of nodes, and wherein the computed parameters comprise interference nulling parameters for a receiving beam to mitigate an angle of arrival of the interference source.

14. The method of claim 1, wherein said analyzing the data on the RF links to compute parameters for the alternate routing paths comprises pre-computing a cost function for the alternate routing paths by the controller, wherein the cost function is used to prioritize the alternate routing paths to be used to route the data packet from the source node to the destination node.

15. The method of claim 14, wherein the cost function takes into account the received statistics on the delivery of the data packets from the source node to the destination node using the alternate routing paths.

17

16. The method of claim 15, wherein said analyzing comprises deriving a probability of delivery of a data packet from the source node to the destination node using the alternate routing paths, and wherein the cost function takes into account the probability of delivery of the data packet.

17. The method of claim 14, wherein the received statistics comprise information on a failed link between a pair of nodes in one or more of the alternate routing paths from the source node to the destination node.

18. The method of claim 14, wherein the received statistics comprise a link quality between a pair of nodes in one or more of the alternate routing paths from the source node to the destination node.

19. The method of claim 14, wherein said adaptively adjusting the parameters for the alternate routing paths comprises re-computing the cost function for the alternate routing paths.

20. The method of claim 14, wherein the cost function prioritizes the alternate routing paths to maximize a probability of delivery of a data packet from the source node to the destination node.

21. The method of claim 20, wherein for the data packet that requires low latency from the source node to the destination node, the alternate routing paths are prioritized to minimize the number of hops from the source node to the destination node.

22. The method of claim 20, wherein one or more copies of the data packet are transmitted over the alternate routing paths from the source node to the destination node.

23. The method of claim 1, wherein the parameters for the alternate routing paths comprise parameters for modulation and error coding between a pair of node from the source node to the destination node for each of the alternate routing paths.

24. The method of claim 23, wherein said adaptively adjusting the parameters for the alternate routing paths comprises adjusting the modulation and error coding for the plurality of node pairs from the source node to the destination for the alternate routing paths as a function of the received statistics on the delivery of the data packets from the source node to the destination node.

25. The method of claim 23, wherein said adaptively adjusting the parameters for the alternate routing paths comprises adjusting the modulation and error coding for the plurality of node pairs from the source node to the destination for the alternate routing paths as a function of a link quality from the source node to the destination node.

26. The method of claim 1, wherein the parameters comprise beam parameters for a directed beam between the plurality of node pairs from the source node to the destination node for the alternate routing paths.

27. The method of claim 1, wherein the parameters for the alternate routing paths comprise information to fragment the data packet into a plurality of data fragments and information for sending the plurality of data fragments over a plurality of alternate routing paths from the source node to the destination node.

28. The method of claim 1, wherein said analyzing comprises determining a pair of unwanted transmit-receive nodes wherein a transmission from the unwanted transmit node interferes with a receiving function at the unwanted receiving node in an interference direction, and wherein the parameters are computed to avoid the interference between the pair of unwanted transmit-receive nodes in the interference direction.

29. The method of claim 1, wherein the wireless data network operates at a single frequency across the plurality of nodes.

18

30. The method of claim 1, further comprising synchronizing transmitting and receiving of data packets across the plurality of nodes of the wireless data network.

31. A system of a multi-hop mesh wireless data network, comprising:

one or more memories configured to store machine instructions; and

a processor configured to execute the machine instructions to:

characterize data on radio frequency (RF) links at a plurality of nodes of the wireless data network;

analyze the data on the RF links to compute parameters for one or more alternate routing paths from a source node to a destination node in one or more hops across the wireless data network;

load the parameters for the alternate routing paths to the source node to direct the source node on how to route a data packet from the source node to the destination node through any intermediate nodes;

receive statistics on the delivery of a plurality of data packets from the source node to the destination node; and

adaptively adjust the parameters for the alternate routing paths in response to the data on the RF links and the statistics on the delivery of data packets.

32. The system of claim 31, wherein the processor is configured to characterize the data by being further configured to control a pair of the plurality of nodes to collect signal strength data on the RF link between the pair of nodes.

33. The system of claim 32, wherein the processor is further configured to command a transmitting node of to transmit a probe packet in a plurality of transmit directions around the transmitting node, and to command a listening node, for each of the transmitting directions of the probe packet, to collect received signal strength data of the probe packet in a plurality of incident directions around the listening node.

34. The system of claim 31, wherein the processor is configured to analyze the data by being further configured to detect one or more multi-path reflectors that reflect an RF signal of an RF link between a pair of the plurality of nodes.

35. The system of claim 34, wherein the computed parameters comprise beam parameters to form a transmission beam and a receiving beam that take advantage of the multi-path reflectors in a non-line-of-sight path between the pair of nodes.

36. The system of claim 31, wherein the processor is configured to analyze the data to compute parameters for the alternate routing paths by being further configured to pre-compute a cost function for the alternate routing paths, wherein the cost function is used to prioritize the alternate routing paths to be used to route the data packet from the source node to the destination node.

37. The system of claim 36, wherein the processor is configured to adaptively adjust the parameters for the alternate routing paths by being further configured to re-compute the cost function for the alternate routing paths.

38. The system of claim 36, wherein the cost function prioritizes the alternate routing paths to maximize a probability of delivery of a data packet from the source node to the destination node.

39. The system of claim 31, wherein the processor is configured to analyze the data by being further configured to determine a pair of unwanted transmit-receive nodes wherein a transmission from the unwanted transmit node interferes with a receiving function at the unwanted receiving node in an interference direction, and wherein the parameters are com-

puted to avoid the interference between the pair of unwanted transmit-receive nodes in the interference direction.

40. The system of claim 31, wherein the wireless data network operates at a single frequency across the plurality of nodes, and wherein the processor is further configured to synchronize transmitting and receiving of data packets across the plurality of nodes of the wireless data network.

* * * * *